	Type	Hits	Search Text	DBs
1	BRS	17	S61 and index	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
2	BRS	14	S62 and sub	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
3	BRS	20	S60 and category	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
4	BRS	948	updat$3 with dynamic$5 with display	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
5	BRS	51	S59 and S58	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
6	BRS	18	S55 and S56	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
7	BRS	2749	"stock quote"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
8	BRS	1777	category with subcategory	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
9	BRS	1841	dynamic$5 near4 index$3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

Your search matched **19** of **1128145** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

'dynamic indexing' Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

1 **A dynamic indexing structure for searching time-series patterns**
Young in Kim; Youngbae Park; Jonghoon Chun;
Computer Software and Applications Conference, 1996. COMPSAC '96., Proceedings of 20th International , 21-23 Aug. 1996
Pages:270 - 275

[Abstract] [PDF Full-Text (420 KB)] **IEEE CNF**

---

2 **Efficient lossless coding scheme for vector quantisation using dynamic index mapping**
Seung Jun Lee; Kyeong Ho Yang; Chul Woo Kim; Choong Woong Lee;
Electronics Letters , Volume: 31 , Issue: 17 , 17 Aug. 1995
Pages:1426 - 1427

[Abstract] [PDF Full-Text (164 KB)] **IEE JNL**

---

3 **A method for dynamic indexing of large image databases**
Hendra, Y.; Gopalan, R.P.; Nair, M.G.;
Systems, Man, and Cybernetics, 1999. IEEE SMC '99 Conference Proceedings. 1999 IEEE International Conference on , Volume: 1 , 12-15 Oct. 1999
Pages:302 - 307 vol.1

[Abstract] [PDF Full-Text (632 KB)] **IEEE CNF**

---

4 **The DC-tree: a fully dynamic index structure for data warehouses**
Ester, M.; Kohlhammer, J.; Kriegel, H.-P.;
Data Engineering, 2000. Proceedings. 16th International Conference on , 29 Feb.-3 March 2000
Pages:379 - 388

[Abstract] [PDF Full-Text (92 KB)] **IEEE CNF**

---

5 **System dynamic index for market power mitigation in the restructuring electricity industry**

Jian Yang; Jordan, G.;
Power Engineering Society Summer Meeting, 2000. IEEE , Volume: 4 , 16-20 July 2000
Pages:2217 - 2222 vol. 4

[Abstract] [PDF Full-Text (524 KB)] IEEE CNF

---

6 **Dynamic index assessment for voltage stability in electric power systems**
Delfino, B.; Denergri, G.B.; Invernizzi, M.; Milano, F.;
Power System Technology, 2000. Proceedings. PowerCon 2000. International Conference on , Volume: 1 , 4-7 Dec. 2000
Pages:397 - 402 vol.1

[Abstract] [PDF Full-Text (536 KB)] IEEE CNF

---

7 **The GC-tree: a high-dimensional index structure for similarity search in image databases**
Guang-Ho Cha; Chin-Wan Chung;
Multimedia, IEEE Transactions on , Volume: 4 , Issue: 2 , June 2002
Pages:235 - 247

[Abstract] [PDF Full-Text (523 KB)] IEEE JNL

---

8 **Motive simulation table mixed sensitivity design applying sampled-data control theory**
Ma Jie; Yao Yu; Fang Qiang;
Advanced Motion Control, 2000. Proceedings. 6th International Workshop on , 30 March-1 April 2000
Pages:211 - 215

[Abstract] [PDF Full-Text (264 KB)] IEEE CNF

---

9 **Heart rate variability characterization using correlation dimension**
Xiaobo Miao; Wei He; Hao Yang; Heng-Ming Tail;
Circuits and Systems, 2002. MWSCAS-2002. The 2002 45th Midwest Symposium on , Volume: 1 , 4-7 Aug. 2002
Pages:I - 447-50 vol.1

[Abstract] [PDF Full-Text (319 KB)] IEEE CNF

---

10 **Massively parallel distributed feature extraction in textual data mining using HDDI™**
Kuntraruk, J.; Pottenger, W.M.;
High Performance Distributed Computing, 2001. Proceedings. 10th IEEE International Symposium on , 7-9 Aug. 2001
Pages:363 - 370

[Abstract] [PDF Full-Text (520 KB)] IEEE CNF

---

11 **Design and implementation of a communication mechanism for WEBGIS based on a distributed component model**
Chaozhen Guo; Qingming Wang; Yuan Xu;
Computer Supported Cooperative Work in Design, The Sixth International Conference on, 2001 , 12-14 July 2001
Pages:558 - 561

[Abstract] [PDF Full-Text (332 KB)] IEEE CNF

12 **Adaptive shaping and filtering of laser beams using photorefractive elements**
Roosen, G.; Godard, A.; Meilhac, L.; Dubreuil, N.; Michel, J.;
Lasers and Electro-Optics, 2001. CLEO '01. Technical Digest. Summaries of papers presented at the Conference on , 6-11 May 2001
Pages:341

[Abstract] [PDF Full-Text (112 KB)] IEEE CNF

13 **The design of the cell tree: an object-oriented index structure for geometric databases**
Gunther, O.;
Data Engineering, 1989. Proceedings. Fifth International Conference on , 6-10 Feb. 1989
Pages:598 - 605

[Abstract] [PDF Full-Text (676 KB)] IEEE CNF

14 **Major-minor loop chips adapted for associative search in relational data base**
Hsu Chang; Nigam, A.;
Magnetics, IEEE Transactions on , Volume: 14 , Issue: 6 , Nov 1978
Pages:1123 - 1128

[Abstract] [PDF Full-Text (664 KB)] IEEE JNL

15 **Changes induced by tilting on Poincare plots and symbolic dynamic patterns of HRV compared to spectral indexes in post-MI and normal subjects**
D'Addio, G.; Pinna, G.D.; Maestri, R.; Acanfora, D.; Ranaudo, E.; Furgi, G.; Rengo, F.;
Computers in Cardiology 2000 , 24-27 Sept. 2000
Pages:449 - 452

[Abstract] [PDF Full-Text (380 KB)] IEEE CNF

1 2 Next

Terms used **display update** Found **56** of **377 searched out of 377.**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 56 Result page: **1** 2 3 next

Relevance scale

1 An incremental access method for ViewCache: concept, algorithms, and cost analysis
Nicholas Roussopoulos
September 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 3
Full text available: pdf(1.71 MB) Additional Information: full citation, abstract, references, citings, index terms, review

A ViewCache is a stored collection of pointers pointing to records of underlying relations needed to materialize a view. This paper presents an Incremental Access Method (IAM) that amortizes the maintenance cost of ViewCaches over a long time period or indefinitely. Amortization is based on deferred and other update propagation strategies. A deferred update strategy allows a ViewCache to remain outdated until a query needs to selectively or ...

Keywords: terms

2 Wave-indices: indexing evolving databases
Narayanan Shivakumar, Héctor García-Molina
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2
Full text available: pdf(1.68 MB) Additional Information: full citation, abstract, references, index terms

In many applications, new data is being generated every day. Often an index of the data of a past window of days is required to answer queries efficiently. For example, in a warehouse one may need an index on the sales records of the last week for efficient data mining, or in a Web service one may provide an index of Netnews articles of the past month. In this paper, we propose a variety of wave indices where the data of a new day can be efficiently added, and old data can ...

3 Window real objects: a distributed shared memory for distributed implementation of GUI applications
Noboru Koshizuka, Ken Sakamura
December 1993 **Proceedings of the 6th annual ACM symposium on User interface software and technology**
Full text available: pdf(1.31 MB) Additional Information: full citation, references, index terms

Keywords: BTRON, distributed shared memory, graphical user interface, multiuser interface, parallel programming, window system

4 Direct spatial search on pictorial databases using packed R-trees
Nick Roussopoulos, Daniel Leifker
May 1985 **ACM SIGMOD Record , Proceedings of the 1985 ACM SIGMOD international conference on Management of data**, Volume 14 Issue 4
Full text available: pdf(1.26 MB) Additional Information: full citation, references, citings, index terms

5 Distributed, Web-based GIS: A systematic approach to reduction of user-perceived response time for GIS web services
Shengru Tu, Xiangfeng He, Xuefeng Li, Jay J. Ratcliff
November 2001 **Proceedings of the 9th ACM international symposium on Advances in geographic information systems**
Full text available: pdf(1.62 MB) Additional Information: full citation, abstract, references, citings, index terms

The research of Internet distribution of GIS contents is still in its infancy. This paper reports an implementation of a systematic approach to optimize Internet distribution of GIS datasets. The goal is to reduce user-perceived response time and improve users' navigation efficiency.On the server side, a large GIS dataset associated with a map is decomposed into small blocks and organized into an overview guiding hierarchy. On the client side, locality-based caching and pre-fetching techniques a ...

Keywords: GIS, Java, applet, caching, hierarchical model, pre-fetching, response time, web

6 Multidimensional access methods
Volker Gaede, Oliver Günther
June 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 2
Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, references, citings, index terms

Search operations in databases require special support at the physical level. This is true for conventional databases as well as spatial databases, where typical search operations include the point query (find all objects that contain a given search point) and the region query (find all objects that overlap a given search region). More than ten years of spatial database research have resulted in a great variety of multidimensional access methods to support ...

Keywords: data structures, multidimensional access methods

7 A hierarchical access control model for video database systems
Elisa Bertino, Jianping Fan, Elena Ferrari, Mohand-Said Hacid, Ahmed K. Elmagarmid, Xingquan Zhu
April 2003 **ACM Transactions on Information Systems (TOIS)**, Volume 21 Issue 2
Full text available: pdf(6.27 MB) Additional Information: full citation, abstract, references, index terms

Content-based video database access control is becoming very important, but it depends on the progresses of the following related research issues: (a) efficient video analysis for supporting semantic visual concept representation; (b) effective video database indexing structure; (c) the development of suitable video database models; and (d) the development of access control models tailored to the characteristics of video data. In this paper, we propose a novel approach to support multilevel acce ...

Keywords: Video database models, access control, indexing schemes

8 The model, language, and implementation of an object-oriented multimedia knowledge base management system

Hiroshi Ishikawa, Fumio Suzuki, Fumihiko Kozakura, Akifumi Makinouchi, Mika Miyagishima, Yoshio Izumida, Masaaki Aoshima, Yasuo Yamane
March 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 1

Full text available: pdf(3.23 MB) Additional Information: full citation, abstract, references, citings, index terms, review

New applications such as CAD, AI, and hypermedia require direct representation and flexible use of complex objects, behavioral knowledge, and multimedia data. To this end, we have devised a knowledge base management system called Jasmine. An object-oriented approach in a programming language also seems promising for use in Jasmine. Jasmine extends the current object-oriented approach and provides the following features. Our object model is based on functional data models and well-establis ...

9 The design of POSTGRES

Michael Stonebraker, Lawrence A. Rowe
June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data**, Volume 15 Issue 2

Full text available: pdf(1.91 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents the preliminary design of a new database management system, called POSTGRES, that is the successor to the INGRES relational database system. The main design goals of the new system are to provide better support for complex objects, provide user extendibility for data types, operators and access methods, provide facilities for active databases (i.e., alerters and triggers) and inferencing including forward- ...

10 Versioning a full-text information retrieval system

Peter G. Anick, Rex A. Flynn
June 1992 **Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval**

Full text available: pdf(1.53 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, we present an approach to the incorporation of object versioning into a distributed full-text information retrieval system. We propose an implementation based on "partially versioned" index sets, arguing that its space overhead and query-time performance make it suitable for full-text IR, with its heavy dependence on inverted indexing. We develop algorithms for computing both historical queries and time range queries and show how these algorithms can be applied to ...

11 Data structures for dynamic queries: an analytical and experimental evaluation

Vinit Jain, Ben Shneiderman
June 1994 **Proceedings of the workshop on Advanced visual interfaces**

Full text available: pdf(2.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Dynamic Queries is a querying technique for doing range search on multi-key data sets. It is a direct manipulation mechanism where the query is formulated using graphical widgets and the results are displayed graphically preferably within 100 milliseconds.This paper evaluates four data structures, the multilist, the grid file, k-d tree and the quad tree used to organize data in high speed storage for dynamic queries. The effect of factors like size, distribution and dimensionalit ...

12 Tools: A translator's workstation

Eugenio Picchi, Carol Peters, Elisabetta Marinai
August 1992 **Proceedings of the 14th conference on Computational linguistics - Volume 3**

Full text available: pdf(389.66 KB) Additional Information: full citation, abstract, references

A description is given of the present state of development of a workstation that has been designed to provide the translator with efficient and easy-to-use computational tools. The

aim is to offer translators fast and flexible on-line access to existing dictionary databases and bilingual text archives and also to supply them with facilities for updating, adding to and personalizing the system data archives with their own material.

13 Special issue on spatial database systems: Management of multidimensional discrete data

Peter Baumann

October 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 4

Full text available: pdf(2.30 MB) Additional Information: full citation, abstract, references, citings

Spatial database management involves two main categories of data: vector and raster data. The former has received a lot of in-depth investigation; the latter still lacks a sound framework. Current DBMSs either regard raster data as pure byte sequences where the DBMS has no knowledge about the underlying semantics, or they do not complement array structures with storage mechanisms suitable for huge arrays, or they are designed as specialized systems with sophisticated imaging functionality, but n ...

Keywords: Multimedia database systems, image database systems, spatial index, tiling

14 Fast image retrieval using color-spatial information

Beng Chin Ooi, Kian-Lee Tan, Tat Seng Chua, Wynne Hsu

May 1998 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 2

Full text available: pdf(496.55 KB) Additional Information: full citation, abstract, index terms

In this paper, we present an image retrieval system that employs both the color and spatial information of images to facilitate the retrieval process. The basic unit used in our technique is a *single-colored cluster*, which bounds a homogeneous region of that color in an image. Two clusters from two images are similar if they are of the same color and overlap in the image space. The number of clusters that can be extracted from an image can be very large, and it affects the accuracy of ret ...

Keywords: Color-spatial information, Content-based retrieval, Sequenced multi-attribute tree, Single-colored cluster

15 Dynamic vp-tree indexing for *n*-nearest neighbor search given pair-wise distances

Ada Wai-chee Fu, Polly Mei-shuen Chan, Yin-Ling Cheung, Yiu Sang Moon

July 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 9 Issue 2

Full text available: pdf(232.09 KB) Additional Information: full citation, abstract, index terms

For some multimedia applications, it has been found that domain objects cannot be represented as feature vectors in a multidimensional space. Instead, pair-wise distances between data objects are the only input. To support content-based retrieval, one approach maps each object to a *k*-dimensional (*k*-d) point and tries to preserve the distances among the points. Then, existing spatial access index methods such as the R-trees and KD-trees can support fast searching on the resulting

***Keywords**: Content-based retrieval, Indexing, Nearest neighbor search, Pair-wise distances, Updating*

16 Middleware for mobility: MiddleWhere: a middleware for location awareness in ubiquitous computing applications

Anand Ranganathan, Jalal Al-Muhtadi, Shiva Chetan, Roy Campbell, M. Dennis Mickunas

October 2004 **Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware**

Full text available: pdf(326.62 KB) Additional Information: full citation, abstract, references

Location awareness significantly enhances the functionality of ubiquitous computing services and applications, and enriches the way they interact with users and resources in the environment. Many different alternative or complementary location sensing technologies are available. However, these technologies give location information in different formats and with different resolution and confidence. In this paper we introduce "MiddleWhere" a distributed middleware infrastructure for location that ...

17 Analysis of predictive spatio-temporal queries
Yufei Tao, Jimeng Sun, Dimitris Papadias
December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4
Full text available: pdf(575.50 KB) Additional Information: full citation, abstract, references, index terms

Given a set of objects S, a spatio-temporal *window query* q retrieves the objects of S that will intersect the window during the (future) interval q_T. A *nearest neighbor query* q retrieves the objects of S closest to q during q_T. Given a threshold d, a spatio-temporal *join* retrieves the pairs of objects from two datasets that will come within distance d from each other during q
Keywords: Database, histogram, nearest distance, selectivity, spatio-temporal

18 Learning subjective relevance to facilitate information access
James R. Chen, Nathalie Mathé
December 1995 **Proceedings of the fourth international conference on Information and knowledge management**
Full text available: pdf(943.28 KB) Additional Information: full citation, references, index terms

19 DB-6 (databases): XML query processing: Virtual cursors for XML joins
Beverly Yang, Marcus Fontoura, Eugene Shekita, Sridhar Rajagopalan, Kevin Beyer
November 2004 **Proceedings of the Thirteenth ACM conference on Information and knowledge management**
Full text available: pdf(371.20 KB) Additional Information: full citation, abstract, references, index terms

Structural joins are a fundamental operation in XML query processing and a large body of work has focused on index-based algorithms for executing them. In this paper, we describe how two well-known index features -- path indices and ancestor information -- can be combined in a novel way to replace one or more of the physical index cursors in a structural join with virtual cursors. The position of a virtual cursor is derived from the path and ancestor information of a physical c ...

Keywords: XML, evaluation, indexing, join operator

20 Energy-performance trade-offs for spatial access methods on memory-resident data
Ning An, Sudhanva Gurumurthi, Anand Sivasubramaniam, Narayanan Vijaykrishnan, Mahmut Kandemir, Mary Jane Irwin
November 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 3
Full text available: pdf(641.55 KB) Additional Information: full citation, abstract, index terms

The proliferation of mobile and pervasive computing devices has brought energy constraints into the limelight. Energy-conscious design is important at all levels of system architecture, and the software has a key role to play in conserving battery energy on these devices. With the increasing popularity of spatial database applications, and their anticipated deployment on mobile devices (such as road atlases and GPS-based applications), it is critical to examine the energy implications of spatial ...

Keywords: Energy optimization, Multidimensional indexing, Resource-constrained computing, Spatial data

Results 1 - 20 of 56 Result page: **1** 2 3 next

Terms used **display update**

Found **56** of **377**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 21 - 40 of 56

Result page: previous 1 **2** 3 next

Relevance scale

21 External memory algorithms and data structures: dealing with **massive data**

Jeffrey Scott Vitter

June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Full text available: pdf(828.46 KB)

Additional Information: full citation, abstract, references, citings, index terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

22 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available: pdf(9.37 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

23 Content-based image retrieval for multimedia databases: Image database retrieval

utilizing affinity relationships
Mei-Ling Shyu, Shu-Ching Chen, Min Chen, Chengcui Zhang, Kanoksri Sarinnapakorn
November 2003 **Proceedings of the 1st ACM international workshop on Multimedia databases**

Full text available: pdf(554.57 KB) Additional Information: full citation, abstract, references, citings, index terms

Recent research effort in Content-Based Image Retrieval (CBIR) focuses on bridging the gap between low-level features and high-level semantic contents of images as this gap has become the bottleneck of CBIR. In this paper, an effective image database retrieval framework using a new mechanism called the Markov Model Mediator (MMM) is presented to meet this demand by taking into consideration not only the low-level image features, but also the high-level concepts learned from the history of user's ...

24 The Quadtree and Related Hierarchical Data Structures
Hanan Samet
June 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 2

Full text available: pdf(4.87 MB) Additional Information: full citation, references, citings, index terms

25 Index-driven similarity search in metric spaces
Gisli R. Hjaltason, Hanan Samet
December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4

Full text available: pdf(650.64 KB) Additional Information: full citation, abstract, references, index terms

Similarity search is a very important operation in multimedia databases and other database applications involving complex objects, and involves finding objects in a data set *S* similar to a query object *q*, based on some similarity measure. In this article, we focus on methods for similarity search that make the general assumption that similarity is represented with a distance metric *d*. Existing methods for handling similarity search in this setting typically fall into one of ...

Keywords: Hiearchical metric data structures, distance-based indexing, nearest neighbor queries, range queries, ranking, similarity searching

26 Locally adaptive dimensionality reduction for indexing large time series databases
Kaushik Chakrabarti, Eamonn Keogh, Sharad Mehrotra, Michael Pazzani
June 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 2

Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, index terms

Similarity search in large time series databases has attracted much research interest recently. It is a difficult problem because of the typically high dimensionality of the data. The most promising solutions involve performing dimensionality reduction on the data, then indexing the reduced data with a multidimensional index structure. Many dimensionality reduction techniques have been proposed, including Singular Value Decomposition (SVD), the Discrete Fourier transform (DFT), and the Discrete ...

Keywords: Dimensionality reduction, indexing, time-series similarity retrieval

27 The Grid File: An Adaptable, Symmetric Multikey File Structure
J. Nievergelt, Hans Hinterberger, Kenneth C. Sevcik
March 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 1

Full text available: pdf(2.35 MB) Additional Information: full citation, abstract, references, citings, index terms

Traditional file structures that provide multikey access to records, for example, inverted files, are extensions of file structures originally designed for single-key access. They

manifest various deficiencies in particular for multikey access to highly dynamic files. We study the dynamic aspects of file structures that treat all keys symmetrically, that is, file structures which avoid the distinction between primary and secondary keys. We start from a bitmap approach and treat the problem ...

28 Incremental clustering for dynamic information processing

Fazli Can

April 1993 **ACM Transactions on Information Systems (TOIS)**, Volume 11 Issue 2

Full text available: pdf(1.49 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Clustering of very large document databases is useful for both searching and browsing. The periodic updating of clusters is required due to the dynamic nature of databases. An algorithm for incremental clustering is introduced. The complexity and cost analysis of the algorithm together with an investigation of its expected behavior are presented. Through empirical testing it is shown that the algorithm achieves cost effectiveness and generates statistically valid clusters that are compatibl ...

Keywords: best-match cluster search, cluster validity, cover coefficient, dynamic information retrieval environment, information retrieval, information retrieval effectiveness, information retrieval efficiency

29 Spatial priority search: an access technique for scaleless maps

Bruno Becker, Hans-Werner Six, Peter Widmayer

April 1991 **ACM SIGMOD Record , Proceedings of the 1991 ACM SIGMOD international conference on Management of data**, Volume 20 Issue 2

Full text available: pdf(1.04 MB) Additional Information: full citation, references, citings, index terms

30 Database session 8: interactive data exploration: iTopN: incremental extraction of the N most visible objects

Linas Bukauskas, Leo Mark, Edward Omiecinski, Michael H. Böhlen

November 2003 **Proceedings of the twelfth international conference on Information and knowledge management**

Full text available: pdf(312.16 KB) Additional Information: full citation, abstract, references, index terms

The visual exploration of large databases calls for a tight coupling of database and visualization systems. Current visualization systems typically fetch all the data and organize it in a scene tree, which is then used to render the visible data. For immersive data explorations, where an observer navigates in a potentially huge data space and explores selected data regions this approach is inadequate. A scalable approach is to make the database system *observer-aware* and exchange the data ...

Keywords: incremental observer relative data extraction, indexing visibility ranges, moving observer, top most visible objects

31 Path sharing and predicate evaluation for high-performance XML filtering

Yanlei Diao, Mehmet Altinel, Michael J. Franklin, Hao Zhang, Peter Fischer

December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4

Full text available: pdf(543.40 KB) Additional Information: full citation, abstract, references, index terms

XML filtering systems aim to provide fast, on-the-fly matching of XML-encoded data to large numbers of query specifications containing constraints on both structure and content. It is now well accepted that approaches using event-based parsing and Finite State Machines (FSMs) can provide the basis for highly scalable structure-oriented XML filtering systems. The XFilter system [Altinel and Franklin 2000] was the first published FSM-based XML filtering approach. XFilter used a separate FSM per pa ...

Keywords: Nondeterministic Finite Automaton, XML filtering, content-based matching, nested path expressions., path sharing, predicate evaluation, structure matching

32 A unifying model of physical databases

D. S. Batory, C. C. Gotlieb

December 1982 **ACM Transactions on Database Systems (TODS)**, Volume 7 Issue 4

Full text available: pdf(1.91 MB) Additional Information: full citation, abstract, references, citings, index terms

A unifying model for the study of database performance is proposed. Applications of the model are shown to relate and extend important work concerning batched searching, transposed files, index selection, dynamic hash-based files, generalized access path structures, differential files, network databases, and multifile query processing.

Keywords: decomposition, linksets, simple files, unifying model

33 Indexing large metric spaces for similarity search queries

Tolga Bozkaya, Meral Ozsoyoglu

September 1999 **ACM Transactions on Database Systems (TODS)**, Volume 24 Issue 3

Full text available: pdf(281.78 KB) Additional Information: full citation, abstract, references, citings, index terms, review

One of the common queries in many database applications is finding approximate matches to a given query item from a collection of data items. For example, given an image database, one may want to retrieve all images that are similar to a given query image. Distance-based index structures are proposed for applications where the distance computations between objects of the data domain are expensive (such as high-dimensional data) and the distance function is metric. In this paper we consider ...

34 A web-based multimedia database for national flag application

Longzhuang Li, Dehu Qi, Xinqi Zhang

October 2004 **Journal of Computing Sciences in Colleges**, Volume 20 Issue 1

Full text available: pdf(577.56 KB) Additional Information: full citation, abstract, references

Frequently in the database management systems (DBMS) class, we teach students the theory and practical skills required to create efficient Web-based database systems. Generally, at Texas A&M University-Corpus Christi and Lamar University, the projects done in the DBMS class are text-based database systems. However, in this paper we describe a project in which students implement a Web-based multimedia database for the national flag application. The development and implementation of the multimedia ...

35 Database session 8: interactive data exploration: Hierarchical graph indexing

James Abello, Yannis Kotidis

November 2003 **Proceedings of the twelfth international conference on Information and knowledge management**

Full text available: pdf(389.96 KB) Additional Information: full citation, abstract, references, index terms

Traffic analysis, in the context of Telecommunications or Internet and Web data, is crucial for large network operations. Data in such networks is often provided as large graphs with hundreds of millions of vertices and edges. We propose efficient techniques for managing such graphs at the storage level in order to facilitate its processing at the interface level (visualization). The methods are based on a hierarchical decomposition of the graph edge set that is inherited from a hierarchical deco ...

Keywords: graph, index, navigation, visualization

36 Reuse of algorithms: still a challenge to object-oriented programming
Karsten Weihe
October 1997 **ACM SIGPLAN Notices , Proceedings of the 12th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications**, Volume 32 Issue 10

Full text available: pdf(2.32 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper is about reusable, efficient implementations of complex algorithms and their integration into software packages. It seems that this problem is not yetwell understood, and that it is not at all clear how object-oriented and other approaches may contribute to a solution. We analyze the problem and try to reduce it to a few key design goals. Moreover, we discuss various existing approaches in light of these goals, and we briefly report experiences with experimental case studies, in which ...

37 Join processing in relational databases
Priti Mishra, Margaret H. Eich
March 1992 **ACM Computing Surveys (CSUR)**, Volume 24 Issue 1

Full text available: pdf(4.42 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most diffidult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples fro ...

Keywords: database machines, distributed processing, join, parallel processing, relational algebra

38 Modeling and comparing change using spatiotemporal helixes
Anthony Stefanidis, Kristin Eickhorst, Peggy Agouris, Panos Partsinevelos
November 2003 **Proceedings of the 11th ACM international symposium on Advances in geographic information systems**

Full text available: pdf(316.32 KB) Additional Information: full citation, abstract, references, index terms

Spatiotemporal helixes are a novel way to model spatiotemporal change. They represent both the movement of an object, as it is expressed by the trajectory of its center, and the changes of its outline. Accordingly they are highly suitable to communicate the evolution of phenomena as they are captured e.g. in sequences of imagery. In this paper we present the spatiotemporal helix model and introduce spatiotemporal similarity metrics for the comparison of helixes. These metrics allow us to compare ...

Keywords: change modeling, geographic information systems, spatiotemporal analysis

39 Mobile computing: Energy efficient exact kNN search in wireless broadcast environments
Bugra Gedik, Aameek Singh, Ling Liu
November 2004 **Proceedings of the 12th annual ACM international workshop on Geographic information systems**

Full text available: pdf(254.39 KB) Additional Information: full citation, abstract, references, index terms

The advances in wireless communication and decreasing costs of mobile devices have enabled users to access desired information at any time. Coupled with positioning technologies like GPS, this opens up an exciting domain of location based services, allowing a mobile user to query for objects based on its current position. Main bottlenecks in such infrastructures are the draining of power of the mobile devices and the limited network bandwidth available. To alleviate these problems, <i>bro ...

Keywords: air indexes, k-nearest neighbor search, wireless data broadcast

40 Paper session I: techniques: Event-based modeling and processing of digital media
Rahul Singh, Zhao Li, Pilho Kim, Derik Pack, Ramesh Jain
June 2004 **Proceedings of the 1st international workshop on Computer vision meets databases**
Full text available: pdf(740.67 KB) Additional Information: full citation, abstract, references

Capture, processing, and assimilation of digital media-based information such as video, images, or audio requires a unified framework within which signal processing techniques and data modeling and retrieval approaches can act and interact. In this paper we present the rudiments of such a framework based on the notion of "events". This framework serves the dual roles of a conceptual data model as well as a prescriptive model that defines the requirements for appropriate signal processing. Amongs ...

Results 21 - 40 of 56 Result page: previous 1 **2** 3 next

Terms used **display update**

Found **56** of **377**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 41 - 56 of 56

Result page: previous 1 2 **3**

Relevance scale

41 Research sessions: path indexing: Accelerating XPath location steps

Torsten Grust

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings, index terms

This work is a proposal for a database index structure that has been specifically designed to support the evaluation of XPath queries. As such, the index is capable to support *all* XPath axes (including ancestor, following, preceding-sibling, descendant-or-self, *etc.*). This feature lets the index stand out among related work on XML indexing structures which had a focus on regular path expressions (which correspond to the XPath axes children and descendant-or-self plus name tests). I ...

42 Designing and accessing scientific digital libraries: On querying geospatial and georeferenced metadata resources in G-portal

Zehua Liu, Ee-Peng Lim, Wee-Keong Ng, Dion H. Goh

May 2003 **Proceedings of the 3rd ACM/IEEE-CS joint conference on Digital libraries**

Full text available: pdf(92.05 KB) Additional Information: full citation, abstract, references, index terms

G-Portal is a web portal system providing a range of digital library services to access geospatial and georeferenced resources on the Web. Among them are the storage and query subsystems that provide a central repository of metadata resources organized under different projects. In GPortal, all metadata resources are represented in XML (Extensible Markup Language) and they are compliant to some resource schemas defined by their creators. The resource schemas are extended versions of a basic resou ...

43 Implementation of data abstraction in the relational database system INGRES

James Ong, Dennis Fogg, Michael Stonebraker

September 1983 **ACM SIGMOD Record**, Volume 14 Issue 1

Full text available: pdf(611.12 KB) Additional Information: full citation, abstract, references

This paper discusses the design and implementation of an *abstract data type (ADT) facility* which was added to the INGRES database manager. Our implementation of ADTs allows a user to register ADTs and ADT operators with the run-time database manager, declare column values of relations to be instances of ADTs, and formulate queries containing references to ADTs and ADT operators. The user view, implementation, performance, and possible extensions to this new facility are described.

44 Accelerating XPath evaluation in any RDBMS

Torsten Grust, Maurice Van Keulen, Jens Teubner

March 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 1

Full text available: pdf(781.01 KB) Additional Information: full citation, abstract, references, index terms

This article is a proposal for a database index structure, the *XPath accelerator*, that has been specifically designed to support the evaluation of XPath path expressions. As such, the index is capable to support *all* XPath axes (including ancestor, following, preceding-sibling, descendant-or-self, etc.). This feature lets the index stand out among related work on XML indexing structures which had a focus on the child and descendant axes only. The index has been designed with a close ...

Keywords: Main-memory databases, XML, XML indexing, XPath

45 Query processing and optimization: Object-relational management of complex geographical objects

Hans-Peter Kriegel, Peter Kunath, Martin Pfeifle, Matthias Renz

November 2004 **Proceedings of the 12th annual ACM international workshop on Geographic information systems**

Full text available: pdf(191.98 KB) Additional Information: full citation, abstract, references, index terms

Modern database applications including computer-aided design, multimedia information systems, medical imaging, molecular biology, or geographical information systems impose new requirements on the effective and efficient management of spatial data. Particular problems arise from the need of high resolutions for large spatial objects and from the design goal to use general purpose database management systems in order to guarantee industrial-strength. In the past two decades, various stand-alon ...

Keywords: data management, object decomposition, object-relational database, spatial

46 Session 11: multimedia analysis and retrieval: VQ-index: an index structure for similarity searching in multimedia databases

Ertem Tuncel, Hakan Ferhatosmanoglu, Kenneth Rose

December 2002 **Proceedings of the tenth ACM international conference on Multimedia**

Full text available: pdf(525.17 KB) Additional Information: full citation, abstract, references, citings

In this paper, we introduce a novel indexing technique based on efficient compression of the feature space for approximate similarity searching in large multimedia databases. Its main novelty is that state-of-the-art tools from the discipline of data compression are adopted to optimize the complexity-performance tradeoff in large data sets. The design procedure optimizes the query access time by jointly accounting for both database distribution and query statistics. We achieve efficient compress ...

Keywords: approximate similarity searching, clustering, indexing, retrieved information reduction, retrieved set reduction, vector quantization

47 FastMap: a fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets

Christos Faloutsos, King-Ip Lin

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international conference on Management of data**, Volume 24 Issue 2

Full text available: pdf(1.21 MB) Additional Information: full citation, abstract, references, citings, index terms

A very promising idea for fast searching in traditional and multimedia databases is to map objects into points in *k*-d space, using *k* feature-extraction functions, provided by a domain expert [25]. Thus, we can subsequently use highly fine-tuned spatial access methods (SAMs), to answer several types of queries, including the 'Query By Example' type (which translates to a range query); the 'all pairs' query (which translates to a spatial join [8]); the nearest-neighbor or best-match ...

48 A high-performance Web-based system design for spatial data accesses
Shu-Ching Chen, Xinran Wang, Naphtali Rishe, Mark Allen Weiss
November 2000 **Proceedings of the 8th ACM international symposium on Advances in geographic information systems**
Full text available: pdf(787.51 KB) Additional Information: full citation, abstract, index terms

With the increasing use of geographical data in real-world applications, Geographic Information Systems (GISs) have recently emerged as a fruitful area for research. Nowadays, a GIS can be combined with World Wide Web (WWW) techniques to provide information to a multitude of users. A high-performance web-based GIS, called TerraFly, has been developed in order to provide web-based GIS accesses to the general public. The design of TerraFly considers three major aspects including system architec ...

Keywords: GIS, internally distributed multithreading, semantic R-tree

49 Distance-based indexing for high-dimensional metric spaces
Tolga Bozkaya, Meral Ozsoyoglu
June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2
Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, citings, index terms

In many database applications, one of the common queries is to find approximate matches to a given query item from a collection of data items. For example, given an image database, one may want to retrieve all images that are similar to a given query image. Distance based index structures are proposed for applications where the data domain is high dimensional, or the distance function used to compute distances between data objects is non-Euclidean. In this paper, we introduce a distance bas ...

50 On a model of indexability and its bounds for range queries
Joseph M. Hellerstein, Elias Koutsoupias, Daniel P. Miranker, Christos H. Papadimitriou, Vasilis Samoladas
January 2002 **Journal of the ACM (JACM)**, Volume 49 Issue 1
Full text available: pdf(190.89 KB) Additional Information: full citation, abstract, references, index terms, review

We develop a theoretical framework to characterize the hardness of indexing data sets on block-access memory devices like hard disks. We define an indexing workload by a data set and a set of potential queries. For a workload, we can construct an indexing scheme, which is a collection of fixed-sized subsets of the data. We identify two measures of efficiency for an indexing scheme on a workload: *storage redundancy, r* (how many times each item in the data set is stored), and *access over ...*

***Keywords**: Database, index, indexability, lower bounds, multidimensional, query, redundancy*

51 Progressive TINs: algorithms and applications
Anil Maheshwari, Pat Morin, Jörg-Rüdiger Sack
November 1997 **Proceedings of the 5th ACM international workshop on Advances in geographic information systems**
Full text available: pdf(799.78 KB) Additional Information: full citation, references, citings, index terms

52 Spatial and nearest-neighbor queries: Hardware acceleration for spatial selections and joins
Chengyu Sun, Divyakant Agrawal, Amr El Abbadi

June 2003 **Proceedings of the 2003 ACM SIGMOD international conference on Management of data**

Full text available: pdf(744.80 KB) Additional Information: full citation, abstract, references, index terms

Spatial database operations are typically performed in two steps. In the *filtering* step, indexes and the minimum bounding rectangles (MBRs) of the objects are used to quickly determine a set of candidate objects, and in the *refinement* step, the actual geometries of the objects are retrieved and compared to the query geometry or each other. Because of the complexity of the computational geometry algorithms involved, the CPU cost of the refinement step is usually the dominant cost of ...

Keywords: hardware acceleration, spatial join, spatial selection

53 Research sessions: implementation techniques: Implementing database operations using SIMD instructions

Jingren Zhou, Kenneth A. Ross

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

Modern CPUs have instructions that allow basic operations to be performed on several data elements in parallel. These instructions are called SIMD instructions, since they apply a single instruction to multiple data elements. SIMD technology was initially built into commodity processors in order to accelerate the performance of multimedia applications. SIMD instructions provide new opportunities for database engine design and implementation. We study various kinds of operations in a database con ...

54 Equi-depth multidimensional histograms

M. Muralikrishna, David J. DeWitt

June 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international conference on Management of data**, Volume 17 Issue 3

Full text available: pdf(1.13 MB) Additional Information: full citation, references, citings, index terms

55 Analysis of object oriented spatial access methods

Christos Faloutsos, Timos Sellis, Nick Roussopoulos

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data**, Volume 16 Issue 3

Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper provides an analysis of R-trees and a variation (R+-trees) that avoids overlapping rectangles in intermediate nodes of the tree. The main contributions of the paper are the following. We provide the first known analysis of R-trees. Although formulas are given for objects in one dimension (line segments), they can be generalized for objects in higher dimensions as well. We show how the transformation of objects to higher dimensions [HINR83] can be effectively ...

56 R-tree implementation using branch-grafting method

Thomas Schreck, Zhengxin Chen

March 2000 **Proceedings of the 2000 ACM symposium on Applied computing**

Full text available: pdf(415.45 KB) Additional Information: full citation, references, index terms

Keywords: R-tree, R* tree, branch grafting heuristic, spatial data structures

Results 41 - 56 of 56 Result page: previous 1 2 **3**